HotMetrics 2008 Session I: Systems Services, and Energy

Scribe: Brian L. Mark, George Mason University

1 Image Management in a Virtualized Data Center

Presenter: Tingxi Tan, University of Calgary
Discussant:

Question: How do you specify the performance? What does the user need to request in this case?
Response: Performance specification may be in terms of CPU cycles, which gives a low-level specification. Users may also want a high-level specification and map it to a VM. Example: Amazon EC2.

Question: How does the scalability problem depend on the number of physical vs. virtual machines? Note that there is a limit on the number of VMs that can be supported on a given physical machine.
Response: The scalability question certainly depends on the number of physical vs. virtual machines. Moreover, there is also a strong dependence on the *type* of VMs requested. For example, some VMs may be more I/O intensive or more CPU intensive. There is a need for performance models to better understand the scalability issue.

Comment: The presentation focused on the boot time of a VM, a one time cost. To try to understand the problem better, other metrics and aspects of the system could considered, e.g., caching.

Comment: It is important to see the problem in a broader context. There are a lot of optimization problems in this problem space. For example, if thousands of images need to be managed, load balancing should be done. There is the issue of assigning VMs to physical machines. Other issues include I/O, CPU, and networking bottlenecks, as well as provisioning issues. Another issue is the question of how to migrate from one setting to another. The problem space is huge.

Response: In addition, if a customer knows in advance, that the data center is is on a VM environment, the sequence of jobs that it can run can be more dynamic. This could also depend on the kind of communications available between service provider and utility provider.

2 Co-designing the Failure Analysis and Monitoring of Large-Scale Systems

Presenter: Abhisek Chandra, University of Minnesota
Discussant:

Comment: The monitoring system needs to be distributed, but the failure analysis is perhaps best done in a centralized manner.

Response: Good point. A major difficulty lies in collecting the right kind of data. Predictability is also a very difficult issue. Some failure events are more predictable than others.

Question: An interesting thing to look at might be the periodicity of failure events. For example, is there any periodicity in the failure times and/or failure durations? Would it useful to use statistics gathering techniques such as histograms, correlations, etc.?

Comment: Scalability is a huge issue in failure analysis. Smaller systems can be monitored in detail But in a large system, data collection and analysis is very difficult. The bottleneck is the human interpretation of results.
Response: Perhaps visualization tools would be helpful in this respect.
Comment: I do not think that visualization will be very helpful in this context. We don't want to do the analysis "manually."
Response: Data mining techniques may be helpful. Alternative models and tools in the data mining space could be used.

Comment: An important capability we would want is the ability to zoom in to the correct subset of important data. The most important system problems are chronic conditions. It is difficult to identify the chronic conditions happening "under the surface."
Response: Anomaly detection techniques may help.
Comment: If the anomaly detection is too sensitive, this may lead to the triggering of too many events. If it is too coarse, the monitoring system may miss the problem entirely.

Question: Co-design. Security issue 3rd party to verify?
Response: There are interesting security problems NOt familiary with security issues

3 Automatic Request Categorization in Internet Service

Presenter: Abhisek Sharma, University of Southern California
Discussant:

Comment: Problem of determining the number of categories is very important. One approach that has been used in the context of multiclass modeling is clustering, e.g., classes are distinguished on the basis of how the requests cluster together. Perhaps the ICA approach with clustering.
Response: Yes, categorization is very important. The paper by Goldzmidt et al. (EC' 01) uses Bayesian networks to do categorization.

Question: Autocorrelation is going to help a lot in this problem. *Response:* Yes, this is a good point.

Question: You have assumed is your presentation that the traffic process is stationary. Is this a good assumption? How can your proposed techniques handle non-stationarity?
Response: In practice there are some things that you can do to relax this assumption.

Question: Benchmark traces may not always reflect reality accurately, so it is important to test solutions using real-system traces.
Response: Agreed. We are trying to obtain real-system traces. from Microsoft, although this may take some time.

Question: In real-system workloads, anomalies can changes or the system can change over time
Response: Our work did not consider prediction. If you could do prediction, you could determine if there is a change in the system. For example, one could do training on the data in order to detect future changes in the system.

Question: How much domain knowledge should be incorporated into these tools vs. purely blind source separation. What can be automated and what cannot?
Response: Domain knowledge can help you, but can also make you get "stuck." One approach could be to consider two or more different features that could point to an anomaly.

4 Fine-Grained Energy Profiling for Power-Aware Application Design

Presenter: Aman Kansal, Microsoft Research
Discussant:

Question: Is the application-level the right level to consider for energy optimization? It may not be efficient. Would a more layered model be more appropriate for web applications, network applications, and other applications?
Response: Everything above the OS level is considered in our approach. We can make run-time libraries more efficient. Certain things only application will know that the OS would not know.

Question: What opportunities exist for compiler design? Is it possible for the compiler to optimize code to be more energy efficient?
Response: Compilers could insert hints in the code to indicate where some operations may be inefficient.

Comment: Layering may make this counter-productive.

Question: There may be also be a drawback for the developer to worry about energy. Designing to minimize power consumption may result in an increase in code size, which could in turn lead to higher energy consumption.
Response: Yes, care must be taken so that if the code changes do not provide sufficient energy gain, the code should not be changed.

Question: Power control is difficult. Is it a good idea to focus on power consumption as opposed to the more traditional issues of resource consumption? In the latter approach, the joules or watts are not calculated explicitly. *Response:* Some things won't be captured by resource consumption E.g., energy during idle states.

Comment: In energy profiling, one should take into account time series analysis and probability distributions.